

On the Hardness of SAT with Community Structure

Nathan Mull, Daniel J. Fremont, and Sanjit A. Seshia

University of California, Berkeley

Abstract. Recent attempts to explain the effectiveness of Boolean satisfiability (SAT) solvers based on conflict-driven clause learning (CDCL) on large industrial benchmarks have focused on the concept of community structure. Specifically, industrial benchmarks have been empirically found to have good community structure, and experiments seem to show a correlation between such structure and the efficiency of CDCL. However, in this paper we establish hardness results suggesting that community structure is not sufficient to explain the success of CDCL in practice. First, we formally characterize a property shared by a wide class of metrics capturing community structure, including “modularity”. Next, we show that the SAT instances with good community structure according to any metric with this property are still NP-hard. Finally, we study a class of random instances generated from the “pseudo-industrial” *community attachment* model of Giráldez-Cru and Levy. We prove that, with high probability, instances from this model that have relatively few communities but are still highly modular require exponentially long resolution proofs and so are hard for CDCL. We also present experimental evidence that our result continues to hold for instances with many more communities. This indicates that actual industrial instances easily solved by CDCL may have some other relevant structure not captured by the community attachment model.

1 Introduction

Over the last 20 years Boolean satisfiability (SAT) solvers have become widely used tools for solving problems in many domains [17, 26]. This is largely the result of the *conflict-driven clause learning* (CDCL) paradigm, introduced in the mid-1990s [19, 3, 18] and much developed since then. This success of SAT solving in practice is perhaps surprising in light of the NP-hardness of SAT, which is widely interpreted to mean that the problem admits no efficient algorithms. This has led to a line of research trying to answer the basic question: *why does CDCL perform so well in practice?* In other words, what is it about industrial SAT instances that allows them to seemingly avoid the worst-case behavior of CDCL?

One possible explanation is that SAT is significantly easier *on average* than in the worst case. For algorithms like CDCL that are based on resolution (which we will discuss in more detail below), this was ruled out by the discovery that random instances require exponentially-long resolution proofs [6]. Of course, industrial instances are generally highly *non-random*, so another possibility is that such instances tend to fall into a tractable class of problems. For example, SAT is known to be fixed-parameter tractable with respect to various natural parameters such as treewidth and clique-width [25]. However, it is unclear whether these parameters are always small in practice. Moreover, if the goal is to analyze the success of CDCL, the existence of *different* algorithms

that take advantage of small (say) treewidth is not relevant: what matters is whether it correlates with CDCL performance, and in fact there is evidence against this [20].

Parameters more relevant to CDCL are the sizes of *backdoors* [27] and *backbones* [21]. In essence, a backdoor is a set of variables which if assigned cause the instance to become solvable by simplification with no further search, while the backbone is the set of variables which can only be assigned one way in every satisfying assignment. Correlations between the sizes of backdoors and backbones and the performance of CDCL have been observed empirically, and some “structured” instances do seem to have small backdoors [15, 13]. Unfortunately, these experiments have all been limited by the computational difficulty of estimating backdoor sizes, and it is unclear whether they are representative of a majority of large industrial benchmarks.

None of these ideas have adequately covered the whole variety of industrial instances, leaving a significant gap between our theoretical understanding of when SAT is easy and the reality of CDCL’s effectiveness in practice. Of course, despite the intuition that industrial instances have some common underlying structure that explains why CDCL is so effective on them, it is probable that no single explanation suffices. There are particular types of industrial instances that are easy for a specific, known reason that does not apply to all other industrial instances [16]. However, it is still worthwhile to seek general explanations covering as many different types of instances as possible.

One recent approach has focused on the concept of *community structure*, as measured by *modularity* [22]. The variables of an instance with “good community structure” (high modularity) can be partitioned into relatively small sets such that few clauses span multiple sets. It has been found that industrial instances exhibit significantly better community structure than random instances [2], and that community structure does empirically correlate with CDCL performance [23, 12]. This makes community structure a plausible candidate for the “hidden structure” underlying the effectiveness of CDCL on industrial benchmarks. In fact, community structure has been used as the basis for a model of random “pseudo-industrial” instances, the *community attachment* model [12]. It has parameters controlling the number of communities and the density of their inter-connections, allowing it to better reflect the properties of industrial benchmarks.

However, as yet there has been little theoretical analysis connecting community structure to CDCL performance. The only relevant work we are aware of is that of Ganian and Szeider [11], who observe that SAT remains NP-hard for highly modular instances. They also give a tractability result for a parameter “h-modularity” inspired by community structure. However, this parameter is significantly different from the usual modularity, there is no evidence that it is small in industrial benchmarks, and the tractability result is via an algorithm completely different from CDCL.

In this paper, we extend the connection between community structure and worst-case complexity, and establish the first theoretical result on the average-case performance of CDCL on modular instances. Specifically, we:

- Define the *polynomial clique metrics* (PCMs), a broad class of graph metrics that includes modularity and other popular measures of graph clustering (Section 3.1).
- Show that the set of SAT instances which have “good community structure” according to any PCM is still NP-hard (Section 3.2).
- Prove that on random unsatisfiable instances from the community attachment model that have fewer than $\Theta(n^{1/10})$ communities but can still be highly modular, CDCL takes exponential time with high probability (Section 4).

- Give experimental evidence that our result continues to hold for instances with $\Theta(n^\alpha)$ communities for $\alpha \gg 1/10$ (Section 5).

Based on these results, we suggest that community structure by itself may not be an adequate explanation for the effectiveness of CDCL in practice. We begin in Section 2 with background on SAT, CDCL, and community structure both generally and as recently applied to SAT, and conclude in Section 5 with a discussion of our results and some directions for future work.

2 Background

2.1 SAT

The *Boolean satisfiability* or SAT problem is to decide, given a Boolean formula $\varphi(\mathbf{x})$ over a vector of variables \mathbf{x} , whether or not there is a *satisfying assignment* to \mathbf{x} that makes the formula true. In this paper, we make the common assumption that the formula φ is in *conjunctive normal form* (CNF): it is a conjunction $\psi_1 \wedge \cdots \wedge \psi_m$ of *clauses*, where each clause ψ_i is a disjunction $\ell_{i1} \vee \cdots \vee \ell_{ik}$. Here each ℓ_{ij} is a *literal*: either a variable from \mathbf{x} or the negation of such a variable. We also assume that every clause has the same length k . A formula φ satisfying these conditions is called a k -SAT formula.

Given a partial assignment ρ to some of the variables \mathbf{x} , the *restriction* $\varphi \upharpoonright_\rho$ of $\varphi(\mathbf{x})$ to ρ is the formula obtained from φ by removing all clauses satisfied by ρ and all literals falsified by ρ . We can apply a restriction to any list of clauses analogously. The *size* of the restriction is the number of variables assigned by ρ .

2.2 Resolution and CDCL

Resolution [24] is a fundamental proof system that underlies modern SAT solving algorithms. It consists of a single rule stating that from clauses $(v \vee \mathbf{w})$ and $(\neg v \vee \mathbf{u})$ that have occurrences of v with opposite polarities, we may infer the clause $(\mathbf{w} \vee \mathbf{u})$. As we will see in a moment, the importance of resolution for our purposes is that in order to establish that a formula φ is unsatisfiable, SAT solvers based on CDCL implicitly construct a *resolution refutation* of φ : a derivation of a contradiction (the empty clause) from φ using the resolution rule. This effectively means that the runtime of such a solver cannot be shorter than the length of the shortest such refutation.

To make this precise we need to define what we mean by CDCL. *Conflict-driven clause learning* [18] describes a class of algorithms that extend the Davis–Putnam–Logemann–Loveland (DPLL) algorithm [7]. DPLL is a classical search algorithm that assigns each variable in turn, backtracking if a clause is falsified by the current assignment. If at any point there is a clause with only a single unassigned variable, then that variable can immediately be given the assignment which satisfies the clause — a rule called *unit propagation*. If we eventually assign every variable, then we have found a satisfying assignment; otherwise, the search will backtrack all the way to the top level, every possible assignment will have been tried, and the formula is unsatisfiable.

CDCL-type algorithms augment this procedure by *learning* at every backtrack point a new clause that summarizes the reason why the current partial assignment falsifies the formula [19, 3]. This *conflict clause* C is derived by resolving the falsified clause F with one or more other clauses that were used to assign variables in F by unit propagation.

As a result C is always derivable from the original formula φ by resolution, and writing out all clauses learned by CDCL when φ is unsatisfiable gives a resolution refutation of φ [4]. So the shortest such refutation gives a lower bound for the runtime of CDCL. This is true regardless of the heuristics used by the particular CDCL variant to decide which variable to assign and its polarity, how exactly to derive the conflict clause, and when to restart search from the beginning (see [4] for a more precise statement). We also note that pre- or inprocessing techniques that add no clauses (e.g. blocked clause elimination [14]) or add only clauses derived via resolution (e.g. variable elimination [8]) will not affect the lower bound.

2.3 Random SAT Instances

To study the performance of CDCL on “typical” instances, we use the framework of *average-case complexity*, which analyzes the efficiency of algorithms on *random* instances drawn from a particular distribution. We will be interested in complexity lower bounds that hold for almost all sufficiently large instances:

Definition 1. An event X occurs **with high probability** in terms of n if $\Pr[X] \rightarrow 1$ as $n \rightarrow \infty$.

For example, if flipping n fair coins, with high probability at least 49% will be heads.

Perhaps the simplest distribution over SAT instances arises from fixing the numbers of variables, clauses, and variables per clause, and then sampling uniformly:

Definition 2. $F_k(n, m)$ is the uniform distribution over k -CNF formulas with n variables and m clauses.

This *random k -SAT model* has been widely studied, and is known to be difficult on average for CDCL (for clause-variable ratios in a certain range) by the resolution lower bound discussed above: with high probability, a random unsatisfiable instance has only exponentially long resolution refutations [5]. As we will discuss shortly, our work extends this result to a more recent random SAT model that favors instances that are “pseudo-industrial” in the sense of having good community structure.

2.4 Community Structure

The notion of community structure has a long history in many fields [10]. The essential idea is that graphs with “good community structure” can be broken into relatively small pieces, *communities*, that are densely connected internally but only sparsely connected to each other. There are a number of metrics which have been proposed to make this notion formal, of which one of the most popular is *modularity* [22]. We consider unweighted graphs as weighted graphs with all weights 1.

Definition 3. Let $G = (V, E)$ and let $\delta = \{C_1, \dots, C_n\}$ be a vertex partition. Let $\deg v$ be the degree of v , and $w(x, y)$ be the weight of the edge (x, y) or zero if there is no such edge. The **modularity** (or *Q-value*) of G is

$$Q = \max_{\delta} \sum_{C \in \delta} \left[\frac{\sum_{x, y \in C} w(x, y)}{\sum_{x, y \in V} w(x, y)} - \left(\frac{\sum_{x \in C} \deg x}{\sum_{x \in V} \deg x} \right)^2 \right].$$

While work on community structure in SAT instances has focused on modularity, there are several competing metrics that have been used to measure community structure in other domains. In Appendix A of this paper, we consider four: silhouette index, conductance, coverage, and performance [1].

Finally, we introduce notation for two graphs that will be useful in this paper: K_n , the complete graph on n vertices, and K_n^m , consisting of m disjoint copies of K_n .

2.5 SAT and Community Structure

Recent work on the community structure of SAT instances begins by associating to each instance its *variable incidence graph* (also known as the *primal graph*).

Definition 4. Let φ be a CNF formula. The **variable incidence graph (VIG)** of φ is the graph $G_\varphi = (V, E)$ where V is the set of all variables occurring in φ and E is the set $\{(v_1, v_2) : v_1, v_2 \in V \text{ and they appear together in some clause of } \varphi\}$.

Some works use a *weighted* version of this graph with $w(v_1, v_2) = \sum_{cl} \left[1 / \binom{|cl|}{2} \right]$, where the sum is over all clauses in which both v_1 and v_2 appear [2, 12]. This ensures that each clause contributes an equal amount to the total weight of the graph regardless of its length. Our results apply to both the weighted and unweighted versions.

Obviously, the graph G_φ does not preserve all information about the instance φ . In particular, the polarities of the literals are ignored. But the graph does capture significant structural information: for example, if the graph has two connected components on variables \mathbf{x} and \mathbf{y} then the formula $\varphi(\mathbf{x}, \mathbf{y})$ can be split into $\psi(\mathbf{x}) \wedge \chi(\mathbf{y})$ and each subformula solved independently. In practice a perfect decomposition is rare, but one into *almost* independent parts is more plausible. This is exactly the idea of community structure, and leads us naturally to consider applying modularity to SAT instances.

Definition 5. The **modularity** of a formula φ is the modularity of G_φ .

As was mentioned earlier, it has been found empirically that modularity correlates with CDCL performance [23]. This is a claim about the *average* behavior of CDCL over a wide variety of industrial benchmarks, not about its behavior on any specific instance. Thus it is naturally formalized in the average-case complexity framework discussed above, by giving a distribution that favors instances that are “industrial” in character. One such proposal, based on the idea that the key commonality of industrial instances is their good community structure, is the *community attachment* model of Giráldez-Cru and Levy [12]. In addition to the numbers of variables and clauses, this model has parameters controlling the number of communities and the (expected) fraction of clauses that lie within a single community instead of spanning multiple communities.

Definition 6. Let N be a set of n variables. A **partition of N into c communities** is a partition $S = \{S_1, \dots, S_c\}$ of N such that $|S_i| = n/c$. A clause is **within a community** if it contains only variables from a single S_i . A **bridge clause** is a clause whose variables are all in different communities.

Definition 7 ([12]). (Community Attachment Model) Let $n, m, c, k \in \mathbb{N}$ and $p \in [0, 1]$ such that c divides n and $2 \leq k \leq c \leq n/k$. Then $F_k(n, m, c, p)$ is the distribution over k -CNF formulas with n variables and m clauses given by the following procedure: first, choose a random partition of n variables into c communities. With probability p ,

choose a clause uniformly among clauses within a community, and otherwise choose uniformly among bridge clauses. Generate m clauses independently in this way.

Remark 1. We define bridge clauses in a way that matches the community attachment model, but as we will discuss below our results also hold for a modified model where a bridge clause is any clause not within a single community.

Like the random k -SAT model $F_k(n, m)$, the model $F_k(n, m, c, p)$ ranges over k -CNF formulas with n variables and m clauses, and each clause is chosen independently of the others. However, in this model the clauses are of two different types: those lying entirely within a community, and those spread across k different communities. The probability p controls how likely a clause is to be of the first type versus the second.

The idea behind this model is that by picking c and p appropriately, one is likely to obtain instances that decompose into loosely-connected communities, as has been observed in actual industrial instances. More precisely, the expected modularity of an instance drawn from $F_k(n, m, c, p)$ is lower bounded by $p - (1/c)$, so that for nontrivial c highly modular instances can be generated by setting p large enough [12]. Furthermore, Giráldez-Cru and Levy find experimentally that high-modularity instances generated with this model are solved more quickly by CDCL than by look-ahead solvers, and the reverse is true for low-modularity instances [12]. This parallels the same observation for industrial instances versus random instances. Thus, they conclude, $F_k(n, m, c, p)$ is a more realistic model of industrial instances than the random k -SAT model $F_k(n, m)$.

3 Worst-Case Hardness

In this section, we propose a simple class of graph metrics that we argue should include most metrics quantifying community structure. We show that modularity is in fact within the class, as are several other popular graph clustering metrics. However, we demonstrate that the set of SAT instances that have “good community structure” according to any metric in the class is NP-hard. Therefore, no such metric can be a guaranteed indicator of the difficulty of a SAT instance.

3.1 A Class of “Modularity-like” Graph Metrics

We begin by formalizing what we mean by a graph metric.

Definition 8. A **graph metric** is a function m from weighted graphs to $[0, 1]$. Given m and any $\epsilon \in [0, 1]$, $\text{SAT}_{m, \epsilon}$ is the class of all SAT instances φ such that $m(G_\varphi) \geq 1 - \epsilon$.

For example, if m is modularity then $\text{SAT}_{m, \epsilon}$ consists of the “high modularity” formulas, where “high” means any modularity above $1 - \epsilon$.

In general we are interested in graph metrics that represent a notion of community structure, assigning larger values to graphs which have such a structure than those that do not. For such a metric m , consider the following property:

Definition 9. A graph metric m is a **polynomial clique metric (PCM)** if for all $\epsilon > 0$, there is a poly-time computable function $c : \mathbb{N} \rightarrow \mathbb{N}$ with at most polynomial growth and some $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, if K is K_n with any positive edge weights then $m(K^{c(n)}) \geq 1 - \epsilon$.

Remark 2. If using the unweighted version of the variable incidence graph, our proofs will work using a relaxed definition that applies only to $K = K_n$ with unit weights.

In essence, the definition states that for any (sufficiently large) size n , at most a polynomial number of copies of K_n are needed to produce a graph that m considers to have “good community structure”. This is a natural property for modularity-like metrics to have, since copies of K_n are in some sense ideal communities: internally connected as much as possible, with no external edges. Of course we would not consider a single copy of K_n to have good community structure, so the definition of a PCM only requires that such structure be obtained for *some* number of copies at most polynomial in n .

Next we demonstrate that the PCMs are a large class including modularity and several other popular clustering metrics. While the other metrics have not been experimentally evaluated in the context of SAT, this still supports our claim that the PCM property is a natural one for metrics of community structure to have. For lack of space, we defer the definitions and analysis of the metrics other than modularity to Appendix A.

Theorem 1. *Modularity is a PCM.*

Proof. Fix any $\epsilon > 0$ and $n \geq 2$. Let K be K_n with arbitrary positive edge weights, and let $G = K^c$. Let δ be the vertex partition that groups two vertices iff they are in the same copy of K . Then since each community is identical, and there are c communities, $\sum_{x,y \in C} w(x,y) / \sum_{x,y \in V} w(x,y) = 1/c$ and $\sum_{x \in C} \deg x / \sum_{x \in V} \deg x = 1/c$ for any $C \in \delta$. Therefore, $Q(G) \geq c(1/c - (1/c)^2) = 1 - 1/c$. Putting $c = 1/\epsilon$, we have $Q(G) \geq 1 - \epsilon$. Since c is $O(1)$ with respect to n , Q is a PCM. \square

Theorem 2. *Silhouette index, conductance, coverage, and performance are PCMs.*

3.2 Hardness of PCM-Modular Instances

Now we show that the SAT instances which have “good community structure” according to a PCM are no easier in the worst case than any other instance. The PCMs thus form a wide class of metrics which cannot be used as a guaranteed indicator of the difficulty of a SAT instance. Our reduction can be viewed as a variation of that suggested by Ganian and Szeider [11] to show NP-hardness in the specific case of modularity.

Theorem 3. *For any PCM m , the class $\text{SAT}_{m,\epsilon}$ is NP-hard for all $\epsilon > 0$.*

Proof. Given a SAT instance ϕ , we will convert it into an equisatisfiable instance of $\text{SAT}_{m,\epsilon}$ in polynomial time. Let V be the set of all variables occurring in ϕ , along with new variables as necessary so that $|V| \geq n_0$. Fixing a variable x not in V , let ψ be the formula obtained by adding to ϕ all clauses of the form $x \vee y \vee z$ with $y, z \in V$. Clearly, the VIG of ψ is K_n with $n = |V| + 1 \geq n_0$. Furthermore, ϕ and ψ are equisatisfiable, since we can simply assert x to satisfy all the new clauses. Now letting χ be the conjunction of $c(n)$ disjoint copies of ψ (i.e. copies with variables renamed so none are common), the variable incidence graph G of χ is $K_n^{c(n)}$ (with some positive weights). By the PCM property, we have $m(G) \geq 1 - \epsilon$, so $\chi \in \text{SAT}_{m,\epsilon}$. Since χ and ψ are clearly equisatisfiable, so are χ and ϕ , and thus this procedure gives a reduction from SAT to $\text{SAT}_{m,\epsilon}$. Finally, the procedure is polynomial-time since $c(n)$ has at most polynomial growth and can be computed in polynomial time. \square

4 Average-Case Hardness

In contrast to the previous section, we now consider the difficulty of modular instances for a particular class of algorithms, namely those like CDCL which prove unsatisfiability by effectively constructing a resolution refutation. While these results are therefore more specific, they are also much more powerful: they show that modular instances are difficult not just in the worst case but also on average.

Our argument is largely based on the resolution lower bound of Beame and Pitassi [5], which can be used to establish the hardness of instances from the random k -SAT model. In order to use that result, we need to show that most instances from the community attachment model have certain *sparsity* properties used by the proof. So our main steps, detailed in Sections 4.1–4.5 below, are as follows:

1. Define a new distribution $\bar{F}_k(n, m, c, p; m')$ over k -CNF formulas that works by taking a *random subformula* of an instance from the random k -SAT model $F_k(n, m')$.
2. Show that this new distribution is in fact identical to the community attachment model $F_k(n, m, c, p)$.
3. Observe that the sparsity properties are inherited by subformulas, so the sparsity result in [5] for the random k -SAT model $F_k(n, m')$ transfers to the community attachment model $F_k(n, m, c, p)$.
4. Adapt the Beame–Pitassi argument [5] to obtain an exponential lower bound on the resolution refutation length.
5. Conclude that CDCL takes exponential time on unsatisfiable formulas from the community attachment model $F_k(n, m, c, p)$ with high probability.

4.1 Defining the New Distribution

We begin by defining our new distribution $\bar{F}_k(n, m, c, p; m')$, which takes an additional parameter m' that we will specify in Section 4.3.

Definition 10. *Let $n, m, c, k, m' \in \mathbb{N}$ and $p \in [0, 1]$ such that $2 \leq k \leq c \leq n/k$. Then $\bar{F}_k(n, m, c, p; m')$ is the distribution over k -CNF formulas with n variables and m clauses defined by Algorithm 1 (which is such a distribution by Lemma 1 below).*

Lemma 1. *For all parameters satisfying the conditions of Definition 10, Algorithm 1 defines a probability distribution over k -CNF formulas with n variables and m clauses.*

Proof. First we must check that $h/b \leq 1$ so that the algorithm is well-defined. We have

$$\frac{h}{b} = \frac{c \binom{n/c}{k}}{\left(\frac{n}{c}\right)^k \binom{c}{k}} \leq \frac{c \left(\frac{n}{c}\right)^k}{k! \left(\frac{n}{c}\right)^k \left(\frac{c}{k}\right)^k} = \frac{k^k}{k! c^{k-1}} \leq \frac{1}{(k-1)!} \leq 1,$$

since we assume $c \geq k$. Algorithm 1 always terminates, returning a formula ψ from either line 14 or 16. In the first case, ψ is a subset of ϕ , which is drawn from $F_k(n, m')$ and so has k -CNF clauses over n variables. Furthermore, the algorithm does not return from line 14 unless ψ has m clauses. In the second case, ψ is drawn from $F_k(n, m, c, p)$, and so again is a k -CNF formula with n variables and m clauses. \square

Algorithm 1 defining the distribution $\overline{F}_k(n, m, c, p; m')$

```

1: choose  $\phi$  from  $F_k(n, m')$ 
2: choose a uniformly random partition of the  $n$  variables into  $c$  communities
3:  $h \leftarrow c \binom{n/c}{k} / \binom{n}{k}$ 
4:  $b \leftarrow (n/c)^k \binom{c}{k} / \binom{n}{k}$ 
5:  $\psi \leftarrow$  the empty formula on  $n$  variables
6: for all clauses  $C$  of  $\phi$  do
7:   with probability  $p$  do
8:     if  $C$  is within a community then
9:       add  $C$  to  $\psi$ 
10:  otherwise do
11:    with probability  $h/b$  do
12:      if  $C$  is a bridge clause then
13:        add  $C$  to  $\psi$ 
14:  if  $|\psi| = m$  then return  $\psi$  ▷ the algorithm “succeeds”
15: choose a fresh  $\psi$  from  $F_k(n, m, c, p)$ 
16: return  $\psi$  ▷ the algorithm “fails”

```

4.2 Comparing the Distribution to the Community Attachment Model

Next we prove that our definition via Algorithm 1 is equivalent to the usual community attachment definition. Since the algorithm adds each clause independently, in essence this amounts to showing that each clause is within a community with probability p .

Lemma 2. *For any $m' \in \mathbb{N}$, the distribution $\overline{F}_k(n, m, c, p; m')$ is identical to the distribution $F_k(n, m, c, p)$.*

Proof. When Algorithm 1 returns a formula ψ from line 16, ψ is drawn from $F_k(n, m, c, p)$, and so the two distributions are trivially identical. So we need only consider the case when the algorithm returns from line 14. Because the algorithm handles each clause of ϕ independently (until m clauses are added), it suffices to show that when a clause is added to ψ , it is within a community with probability p and is otherwise a bridge clause. Starting from line 7 of Algorithm 1, let C_{comm} be the event that the clause C is within a community, C_{bridge} the event that C is a bridge clause, and C_{added} the event that C is added to ψ . Let A be the event that the algorithm takes the random branch on line 7 instead of the branch on line 10. Then we have

$$\Pr[C_{\text{comm}} | C_{\text{added}}] = \Pr[C_{\text{comm}} | A, C_{\text{added}}] \Pr[A | C_{\text{added}}] + \Pr[C_{\text{comm}} | \overline{A}, C_{\text{added}}] \Pr[\overline{A} | C_{\text{added}}].$$

The second term is zero because \overline{A} means the algorithm takes the branch on line 10 and thus only adds the clause if it is a bridge clause. Likewise, $\Pr[C_{\text{comm}} | A, C_{\text{added}}] = 1$ because the branch on line 7 only adds the clause if it is within a community. So

$$\Pr[C_{\text{comm}} | C_{\text{added}}] = \Pr[A | C_{\text{added}}] = \Pr[C_{\text{added}} | A] \Pr[A] / \Pr[C_{\text{added}}].$$

By straightforward counting arguments, $\Pr[C_{\text{comm}}] = c \binom{n/c}{k} / \binom{n}{k} = h$ and $\Pr[C_{\text{bridge}}] = (n/c)^k \binom{c}{k} / \binom{n}{k} = b$. Since the coin flips on lines 7 and 11 are independent of C , we

have $\Pr[C_{\text{added}}|A] = \Pr[C_{\text{comm}}] = h$ and $\Pr[C_{\text{added}}|\bar{A}] = (h/b) \Pr[C_{\text{bridge}}] = h$. Also $\Pr[A] = p$, so

$$\Pr[C_{\text{added}}] = \Pr[C_{\text{added}}|A] \Pr[A] + \Pr[C_{\text{added}}|\bar{A}] \Pr[\bar{A}] = hp + h(1-p) = h.$$

Plugging these into the expression above we obtain $\Pr[C_{\text{comm}}|C_{\text{added}}] = p$. So each clause added to ψ is within a community with probability p , and otherwise by construction it must be a bridge clause. Therefore when Algorithm 1 returns from line 14, it is equivalent to generating m clauses independently, each of which is a uniformly random clause within a community with probability p , and otherwise a uniformly random bridge clause. So $\bar{F}_k(n, m, c, p; m')$ is identical to $F_k(n, m, c, p)$. \square

4.3 Transferring Subformula-Inherited Properties

Algorithm 1 can “fail” by adding fewer than the desired number of clauses m to ψ , then falling back on the community attachment model as a backup. Otherwise, the algorithm “succeeds”, returning on line 14 a formula that was built up from clauses of ϕ and is therefore a subformula of it. Since our goal is to have the formulas from this distribution inherit properties from ϕ , we need to ensure that Algorithm 1 succeeds with high probability. We can do this by taking m' , the number of clauses in ϕ , to be large enough: then even if a given clause is only added to ψ with a small probability, overall we are likely to add m of them. As we will see in the proof, the probability of adding a clause is roughly $1/c^{k-1}$, so taking m' to be slightly larger than $c^{k-1}m$ will suffice. We use the following standard tail bound.

Lemma 3. *If $B(n, p)$ is the number of successes in n Bernoulli trials each with success probability p , then for $k < pn$ we have*

$$\Pr[B(n, p) \leq k] \leq \exp\left(\frac{-(pn - k)^2}{2pn}\right).$$

Lemma 4. *Suppose that c is $o(n)$, $m \rightarrow \infty$ as $n \rightarrow \infty$, and $m' = (1 + \epsilon)c^{k-1}m$ for some $\epsilon > 0$. Then Algorithm 1 returns from line 14 with high probability.*

Proof. As shown in Lemma 2, the probability that starting from line 7 the clause C will be added to ψ is h . So the probability that Algorithm 1 returns from line 14 is $\Pr[B(m', h) \geq m] = 1 - \Pr[B(m', h) \leq m - 1]$. Now observe that

$$hc^{k-1} = \frac{c^k \binom{n/c}{k}}{\binom{n}{k}} = \frac{n(n-c) \cdots (n-c(k+1))}{n(n-1) \cdots (n-k+1)} \leq 1.$$

Furthermore, we have

$$\lim_{n \rightarrow \infty} hc^{k-1} = \lim_{n \rightarrow \infty} \frac{c^k \binom{n/c}{k}}{\binom{n}{k}} = \lim_{n \rightarrow \infty} \left[\frac{\binom{n/c}{k}}{\frac{(n/c)^k}{k!}} \cdot \frac{n^k}{\binom{n}{k}} \right] = \left[\lim_{n \rightarrow \infty} \frac{\binom{n/c}{k}}{\frac{(n/c)^k}{k!}} \right] \left[\lim_{n \rightarrow \infty} \frac{n^k}{\binom{n}{k}} \right] = 1,$$

where in evaluating the second-to-last limit we use the fact that c is $o(n)$ and so $\lim_{n \rightarrow \infty} (n/c) = \infty$. So for sufficiently large n we have $hc^{k-1} \geq 1 - \epsilon/2(1 + \epsilon)$, and therefore

$$hm' = h(1 + \epsilon)c^{k-1}m \geq \left(1 - \frac{\epsilon}{2(1 + \epsilon)}\right)(1 + \epsilon)m = (1 + \epsilon/2)m.$$

Applying Lemma 3, we have

$$\begin{aligned} \Pr[B(m', h) \leq m - 1] &\leq \exp\left(\frac{-[hm' - (m - 1)]^2}{2hm'}\right) \leq \exp\left(\frac{-[(1 + \epsilon/2)m - m]^2}{2h(1 + \epsilon)c^{k-1}m}\right) \\ &= \exp\left(\frac{-m(\epsilon/2)^2}{2(1 + \epsilon) \cdot hc^{k-1}}\right) \leq \exp\left(\frac{-m\epsilon^2}{8(1 + \epsilon)}\right), \end{aligned}$$

which goes to zero as $m \rightarrow \infty$, and therefore as $n \rightarrow \infty$. So with high probability, Algorithm 1 will return from line 14. \square

Now it is simple to show that subformula-inherited properties are indeed passed down from random k -SAT instances to instances drawn from our distribution. Here “subformula-inherited” simply means that if φ has the property, then any formula made up of a subset of the clauses of φ also has the property. For example, being satisfiable is subformula-inherited, but being unsatisfiable is not.

Lemma 5. *Suppose that c is $o(n)$, $m \rightarrow \infty$ as $n \rightarrow \infty$, $m' = (1 + \epsilon)c^{k-1}m$ for some $\epsilon > 0$, and P is a subformula-inherited property. Then if a formula drawn from $F_k(n, m')$ has property P with high probability, a formula drawn from $\bar{F}_k(n, m, c, p; m')$ has property P with high probability.*

Proof. Run Algorithm 1 to sample from $\bar{F}_k(n, m, c, p; m')$. Let P_ψ and P_ϕ respectively be the events that the returned formula ψ and the formula ϕ from line 1 have property P . Also let R be the event that the algorithm returns from line 14. When the algorithm returns from line 14, ψ is a subformula of ϕ , and since P is inherited by subformulas we have $\Pr[P_\psi | R] \geq \Pr[P_\phi]$. Now as ϕ is drawn from $F_k(n, m')$, the event P_ϕ occurs with high probability, and so $\Pr[P_\psi | R] \rightarrow 1$ as $n \rightarrow \infty$. By Lemma 4, the event R also happens with high probability, so $\Pr[P_\psi] \geq \Pr[P_\psi \wedge R] = \Pr[P_\psi | R] \cdot \Pr[R] \rightarrow 1$ as $n \rightarrow \infty$. Therefore ψ has property P with high probability. \square

Together, Lemmas 2 and 5 show that subformula-inherited properties of random k -SAT instances are also possessed (with high probability) by instances from the community attachment model.

4.4 Proving the Resolution Lower Bounds

Now we transition to adapting the argument of Beame and Pitassi [5]. The proof uses two types of sparsity conditions. Both view a clause C as a set of variables, so that another set of variables X “contains” C if and only if every variable in C is in X .

Definition 11. *A formula is n' -sparse if every set of $s \leq n'$ variables contains at most s clauses.*

Definition 12. *Let $n' < n''$. A formula is (n', n'', y) -sparse if every set of s variables with $n' < s \leq n''$ contains at most ys clauses.*

These are both clearly subformula-inherited.

The Beame–Pitassi argument [5] is broken into three major lemmas, each of which we will use without change. The last lemma establishes the sparsity properties above for the random k -SAT model.

Lemma 6 ([5]). Let $n' \leq n$ and F be an unsatisfiable CNF formula in n variables with clauses of size at most k that is both n' -sparse and $(n'(k+\epsilon)/4, n'(k+\epsilon)/2, 2/(k+\epsilon))$ -sparse. Then any resolution proof P of the unsatisfiability of F must include a clause of length at least $\epsilon n'/2$.

Lemma 7 ([5]). Let P be a resolution refutation of F of size S . Given $\beta > 0$, say the **large clauses** of P are those clauses mentioning more than βn distinct variables. Then with probability at least $1 - 2^{1-\beta t/4} S$, a random restriction of size t sets all large clauses in P to 1.

Lemma 8 ([5]). Let $x > 0$, $1 \geq y > 1/(k-1)$, and $z \geq 4$. Fix a restriction ρ on $t \leq \min\{xn/2, x^{1-1/y(k-1)}n^{1-1/(k-1)}/z\}$ variables. Drawing F from $F_k(n, m)$ with

$$m \leq \frac{y}{e^{1+1/y} 2^{k+1/y}} x^{1/y-(k-1)} n,$$

then with probability at least $1 - 2^{-t} - (2^k + 1)/z^{k-1}$, $F|_{\rho}$ is both $(xn/2, xn, y)$ -sparse and xn -sparse.

We can now combine these to prove the analog of the main theorem of Beame and Pitassi for modular instances. Our argument is almost identical to theirs: the only difference is that we apply Lemma 8 to larger instances from the random k -SAT model,¹ so that our results above will give us sparsity for modular instances of the correct size embedded in them as subformulas.

Theorem 4. Let $k \geq 3$, $0 < \epsilon < 1$, and x, t, z, c be functions of n such that $x > 0$, t and z are $\omega(1)$, c is $o(n)$, and t satisfies the conditions of Lemma 8 for all sufficiently large n . Then with high probability, an unsatisfiable formula drawn from $F_k(n, m, c, p)$ with

$$m \leq \frac{1}{2^{7k/2}(1+\epsilon)c^{k-1}} x^{-(k-2-\epsilon)/2} n$$

does not have a resolution refutation of size $\leq 2^{\frac{\epsilon}{4(k+\epsilon)}xt}/8$.

Proof. Let $S = 2^{\frac{\epsilon}{4(k+\epsilon)}xt}/8$ and let U be the set of unsatisfiable k -CNF formulas with n variables and m clauses. For each $\varphi \in U$ fix a shortest resolution refutation P_{φ} , and let $W \subseteq U$ be the set of φ such that $|P_{\varphi}| \leq S$. Let R be the set of all restrictions of size t , and for any formula φ and $\rho \in R$ let $L(\varphi, \rho)$ be the indicator function for the event that either φ is satisfiable or $P_{\varphi}|_{\rho}$ contains a clause of length at least $\epsilon xn/(k+\epsilon)$. Now for any $\varphi \in W$, by Lemma 7 with $\beta = \epsilon x/(k+\epsilon)$ we have

$$\sum_{\rho} \frac{L(\varphi, \rho)}{|R|} \leq 2^{1-\frac{\epsilon}{4(k+\epsilon)}xt} S = 2^{1-\frac{\epsilon}{4(k+\epsilon)}xt} (2^{\frac{\epsilon}{4(k+\epsilon)}xt}/8) = 1/4.$$

Let X be a random variable defined over a restriction ρ and equal to $\Pr_{\varphi}[L(\varphi, \rho) | \varphi \in W]$, where φ is distributed as $F_k(n, m, c, p)$. Putting a uniform distribution on ρ and

¹ Note that as required by its statement, we are applying Lemma 8 to formulas drawn from $F_k(n, m)$, not to formulas drawn from $\overline{F}_k(n, m, c, p; m')$. Lemmas 6 and 7 work for any formula, so we may use all three lemmas precisely as proved in [5].

writing $q(\psi)$ for the conditional distribution $\Pr_\varphi[\varphi = \psi | \varphi \in W]$,

$$\mathbb{E}_\rho[X] = \sum_\rho \frac{1}{|R|} \Pr_\varphi[L(\varphi, \rho) | \varphi \in W] = \sum_{\psi \in W} q(\psi) \left[\sum_\rho \frac{L(\psi, \rho)}{|R|} \right] \leq \sum_{\psi \in W} q(\psi) \frac{1}{4} = \frac{1}{4}.$$

So by Markov's inequality,

$$\Pr_\rho[X \geq 1/2] \leq \frac{\mathbb{E}_\rho[X]}{1/2} \leq 1/2,$$

and therefore there is some ρ' such that $\Pr_\varphi[L(\varphi, \rho') | \varphi \in W] \leq 1/2$. In other words, there is a restriction that eliminates large clauses from a random $\varphi \in W$ with probability at least $1/2$.

Now let $y = 2/(k + \epsilon)$. Since $k \geq 3$ and $\epsilon < 1$ we have $y \geq 1/(k - 1)$ and

$$\begin{aligned} \frac{y}{e^{1+1/y} 2^{k+1/y}} &= 2 \left[(k + \epsilon) e^{1 + \frac{k+\epsilon}{2}} 2^{k + \frac{k+\epsilon}{2}} \right]^{-1} \geq 2(k + \epsilon)^{-1} e^{-\frac{k}{2} - \frac{3}{2}} 2^{-\frac{3k}{2} - \frac{1}{2}} \\ &= 2(k + \epsilon)^{-1} e^{-3/2} 2^{-1/2} 2^{-k(3 + \log_2 e)/2} \\ &\geq 2(k + 1)^{-1} e^{-3/2} 2^{-1/2} 2^{-2.23k} \geq 2^{-1.23k} 2^{-2.23k} \geq 2^{-7k/2}. \end{aligned}$$

By our assumption on m ,

$$\begin{aligned} (1 + \epsilon) c^{k-1} m &\leq 2^{-7k/2} x^{-(k-2-\epsilon)/2} n = 2^{-7k/2} x^{1/y - (k-1)} n \\ &\leq \frac{y}{e^{1+1/y} 2^{k+1/y}} x^{1/y - (k-1)} n. \end{aligned}$$

Finally, since z is $\omega(1)$ we have $z \geq 4$ for sufficiently large n , and then all the conditions of Lemma 8 are satisfied by y , z , t , and $m' = (1 + \epsilon) c^{k-1} m$. Therefore for a formula φ drawn from $F_k(n, m')$, $\varphi \upharpoonright_{\rho'}$ is simultaneously $(xn/2, xn, 2/(k + \epsilon))$ -sparse and xn -sparse with probability at least $1 - 2^{-t} - (2^k + 1)/z^{k-1}$. Since t and z are $\omega(1)$, φ has this property with high probability. Furthermore, the property is inherited by subformulas, so by Lemma 5 it also holds with high probability for formulas drawn from $\overline{F}_k(n, m, c, p; m')$. Then by Lemma 2 the same is true for formulas drawn from $F_k(n, m, c, p)$.

Now let $n' = 2xn/(k + \epsilon)$. Since $k + \epsilon \geq 3$, we have $n' \leq xn$ and so xn -sparsity implies n' -sparsity. Also note that

$$\frac{xn}{2} = \frac{2xn(k + \epsilon)}{4(k + \epsilon)} = \frac{n'(k + \epsilon)}{4} \quad \text{and} \quad xn = \frac{n'(k + \epsilon)}{2}.$$

So by Lemma 6, when drawing an unsatisfiable formula φ from $F_k(n, m, c, p)$, with high probability every resolution refutation of $\varphi \upharpoonright_{\rho'}$ has a clause of length at least $\epsilon n'/2 = \epsilon xn/(k + \epsilon)$. That is, $\Pr_\varphi[L(\varphi, \rho') | \varphi \in U] \rightarrow 1$ as $n \rightarrow \infty$. So

$$\Pr_\varphi[\varphi \in W | \varphi \in U] = \frac{\Pr_\varphi[\varphi \in W \wedge \overline{L(\varphi, \rho')} | \varphi \in U]}{\Pr_\varphi[\overline{L(\varphi, \rho')} | \varphi \in U]} \leq \frac{\Pr_\varphi[\overline{L(\varphi, \rho')} | \varphi \in U]}{1/2} \rightarrow 0$$

as $n \rightarrow \infty$. Therefore with high probability, an unsatisfiable instance drawn from $F_k(n, m, c, p)$ does not have a resolution refutation of size $\leq S$. \square

Next we instantiate this general result to obtain exponential lower bounds for the refutation length when the number of communities is not too large. We use slightly different arguments for $k \geq 4$ and $k = 3$, again following Beame and Pitassi [5]. As the computations are uninteresting, we defer the proofs to Appendix B. The basic idea is to let x go to zero fast enough that the bound on m required by Theorem 4 is satisfied when $m = O(n)$, but slowly enough that the length bound is of the form $2^{O(n^\lambda)}$.

Theorem 5. *Suppose that $k \geq 4$, $m = O(n)$, and $c = O(n^\alpha)$ for some $\alpha < \frac{k-2}{4(k-1)}$. Then there is some $\lambda > 0$ so that with high probability, an unsatisfiable formula drawn from $F_k(n, m, c, p)$ does not have a resolution refutation of size $2^{O(n^\lambda)}$.*

Theorem 6. *Suppose $m = O(n)$ and $c = O(n^\alpha)$ for some $\alpha < 1/10$. Then there is some $\lambda > 0$ so that with high probability, an unsatisfiable formula drawn from $F_3(n, m, c, p)$ does not have a resolution refutation of size $2^{O(n^\lambda)}$.*

4.5 Deducing a Lower Bound on CDCL Runtime

Finally, we can conclude that unsatisfiable random instances from $F_k(n, m, c, p)$ with sufficiently few communities usually take exponential time for CDCL to solve.

Theorem 7. *If $m = O(n)$ and $c = O(n^\alpha)$ for any $\alpha < 1/10$, the runtime of CDCL on an unsatisfiable formula φ from $F_k(n, m, c, p)$ is exponential with high probability.*

Proof. If φ is unsatisfiable, the runtime of CDCL on φ is lower bounded (up to a polynomial factor) by the length of the shortest resolution refutation of φ [4]. If $k = 3$, then the shortest refutation of φ is exponentially long with high probability by Theorem 6. If instead $k \geq 4$, the same is true by Theorem 5, since $1/10 < \frac{k-2}{4(k-1)}$. Therefore with high probability, CDCL will take exponential time to prove φ unsatisfiable. \square

Remark 3. By picking a sufficiently high clause-variable ratio, we can ensure φ is unsatisfiable with high probability, so that CDCL takes exponential time on average for formulas drawn from $F_k(n, m, c, p)$ (not just the unsatisfiable ones).

We also note that our proof technique is not sensitive to the details of how the community attachment model is defined. For example, changing the definition of a bridge clause so that the variables do not all have to be in different communities requires only minor changes to the proof (detailed in Appendix B).

Theorem 8. *Let $\tilde{F}_k(n, m, c, p)$ be the community attachment model modified so that any clause that is not within a single community counts as a bridge clause. Then if $m = O(n)$ and $c = O(n^\alpha)$ for any $\alpha < 1/10$, the runtime of CDCL on an unsatisfiable formula φ from $\tilde{F}_k(n, m, c, p)$ is exponential with high probability.*

Thus we have showed that similarly to unsatisfiable random k -SAT instances, unsatisfiable random *modular* instances (as formalized by the community attachment model) are hard on average for CDCL as long as they do not have too many communities.

5 Discussion

We have introduced a broad class of “modularity-like” graph metrics, the polynomial clique metrics, and showed that no PCM can be a guaranteed indicator of whether a SAT instance is easy (unless $P = NP$). This is perhaps not too surprising in light of the fact that the VIG throws away the Boolean information in the formula. While the VIG has received the most attention in recent work on community structure, it would be worthwhile to investigate other graph encodings that preserve more information. Regardless, our result does indicate that it may be difficult to define a tractable class of SAT instances based purely on modularity or its variants. Furthermore, by setting up a concrete barrier (the PCM property) that must be avoided to obtain a tractable class, our result can help guide future attempts to find a graph metric that does work.

Our result on the community attachment model $F_k(n, m, c, p)$ is more interesting, as it shows that instances from this model are exponentially hard for CDCL even on average (when c is small enough). An important point is that the result is actually non-trivial when $p < 1$, unlike for $p = 1$. In the latter case there are no bridge clauses, so the instances consist of c independent problems of size n/c , and since we assume $c = O(n^{1/10})$ each problem has size $\Omega(n^{9/10})$. So by the old results on random k -SAT CDCL would take exponential time to solve even the easiest problem, and so likewise for the original instance (with a slightly smaller exponent on n). On the other hand, when $p < 1$ it is conceivable that the bridge clauses could actually make the instances easier, by adding some extra propagation power or easier-to-find contradictions that would make the whole instance easier to solve than any individual community. Our result effectively says that this happens with vanishing probability as $n \rightarrow \infty$.

The case $p = 1$ also brings out an important caveat when interpreting our result as evidence that community structure doesn’t explain CDCL’s effectiveness on industrial instances. Our result shows that such structure isn’t enough to bring random formulas down from exponential-time-on-average to polynomial-time-on-average. However, it could decrease the time from (say) $2^{n^{1/2}}$ to $2^{n^{1/4}}$, which could be the difference between intractability and tractability if n is small enough. On the other hand, given the enormous size of many industrial instances it isn’t clear whether this is really all that is happening. It would be interesting to do experiments on parametrized families of industrial instances to see whether CDCL actually avoids exponential behavior, or if the point of blow-up is just pushed out far enough that we tend not to encounter it in practice.

Another important aspect of our result is the limit on the number of communities. It does not apply when communities have logarithmic size, for example, so that $c = \Theta(n/\log n)$. In fact it is easy to see that the result cannot hold in this case: if one of the communities is unsatisfiable then it will have a polynomial-length resolution refutation, and as $c \rightarrow \infty$ the probability that at least one community is unsatisfiable by itself goes to 1. So with high probability the entire instance has a short refutation, and CDCL could in theory solve it in polynomial time. A clear direction for future work is to see whether improved proof techniques can extend our results to larger numbers of communities, closing the gap between $O(n^{1/10})$ and $\Omega(n/\log n)$. This is also another way our results can inform future experiments: it would be interesting to explore a variety of growth rates for c above $n^{1/10}$ and see how the performance of CDCL changes.

We have done some preliminary experiments along these lines, sampling instances from $F_3(n, 5n, 5n^\alpha, 0.9)$ for a variety of values of n and solving them with MiniSat 2.2.0 [9]. For each value of n we generated 10 instances using the generator from [12]

and averaged their runtimes. Note that every instance had at least 10 communities, so that the expected modularity was at least $p - (1/c) \geq 0.9 - 0.1 = 0.8$. In Figure 1, we plot the results as a function of the community size $n/(5n^\alpha) = n^{1-\alpha}/5$. It




Fig. 1. Average MiniSat runtimes for instances from $F_3(n, 5n, 5n^\alpha, 0.9)$ for $\alpha = 0.1$ (orange), 0.2 (blue), 0.3 (green), and 0.4 (red), plotted as a function of community size ($n^{1-\alpha}/5$).

is clear from the right half of the graph² that the runtime blows up exponentially in the community size (and thus in n) even for α significantly larger than $1/10$. This suggests that improving our result to larger α is likely possible. However, it is important to point out that although all values of α are undergoing exponential growth, the lines for the different values of α are quite close together, indicating that community size is a much more important factor in determining runtime than the total formula size. For example, at a community size of around 270 the $\alpha = 0.1$ instances have 3,000 variables total, while the $\alpha = 0.4$ instances have 165,000. So the latter are more than 50 times larger than the former, but their runtimes are only about 3 times longer. This shows that while larger values of α do not avoid exponential blowup, they can significantly aid performance.

In total, our results indicate that high modularity alone may not be adequate to ensure good performance even on average, but that it could be rewarding to investigate

² On the left half the growth is much slower and close to linear, but since this occurs only for runtimes on the order of a second or less it may be that parsing the formula and initializing the solver dominate the time needed for the actual search.

more refined notions of “good community structure” that somehow restrict the number of communities.

Acknowledgments. The authors thank Vijay Ganesh, Holger Hoos, Zack Newsham, Markus Rabe, Stefan Szeider, and several anonymous reviewers for helpful discussions and comments. This work is supported in part by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1106400, by the Hellman Family Faculty Fund, by gifts from Microsoft and Toyota, and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

References

1. Almeida, H., Guedes, D., Meira Jr, W., Zaki, M.J.: Is there a best quality metric for graph clusters? In: Machine Learning and Knowledge Discovery in Databases, pp. 44–59. Springer (2011)
2. Ansótegui, C., Giráldez-Cru, J., Levy, J.: The community structure of SAT formulas. In: Theory and Applications of Satisfiability Testing - SAT 2012 - 15th International Conference. pp. 410–423 (2012)
3. Bayardo Jr., R.J., Schrag, R.: Using CSP look-back techniques to solve real-world SAT instances. In: Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Innovative Applications of Artificial Intelligence Conference. pp. 203–208 (1997)
4. Beame, P., Kautz, H.A., Sabharwal, A.: Understanding the power of clause learning. In: IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence. pp. 1194–1201 (2003)
5. Beame, P., Pitassi, T.: Simplified and improved resolution lower bounds. In: Foundations of Computer Science, 1996. Proceedings., 37th Annual Symposium on. pp. 274–282. IEEE (1996)
6. Chvátal, V., Szemerédi, E.: Many hard examples for resolution. J. ACM 35(4), 759–768 (1988)
7. Davis, M., Logemann, G., Loveland, D.: A machine program for theorem-proving. Communications of the ACM 5(7), 394–397 (1962)
8. Eén, N., Biere, A.: Effective preprocessing in SAT through variable and clause elimination. In: Theory and Applications of Satisfiability Testing, 8th International Conference. pp. 61–75 (2005)
9. Eén, N., Sörensson, N.: MiniSat, <http://minisat.se/>
10. Fortunato, S.: Community detection in graphs. Physics Reports 486(3-5), 75–174 (2010)
11. Ganian, R., Szeider, S.: Community structure inspired algorithms for SAT and #SAT. In: Theory and Applications of Satisfiability Testing–SAT 2015, pp. 223–237. Springer (2015)
12. Giráldez-Cru, J., Levy, J.: A modularity-based random SAT instances generator. In: 24th Int. Joint Conf. on Artificial Intelligence, IJCAI’15 (2015)
13. Gregory, P., Fox, M., Long, D.: A new empirical study of weak backdoors. In: Principles and Practice of Constraint Programming, 14th International Conference, CP. pp. 618–623 (2008)
14. Jarvisalo, M., Biere, A., Heule, M.: Blocked clause elimination. In: Tools and Algorithms for the Construction and Analysis of Systems, 16th International Conference. pp. 129–144 (2010)
15. Kilby, P., Slaney, J.K., Thiébaux, S., Walsh, T.: Backbones and backdoors in satisfiability. In: Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference. pp. 1368–1373 (2005)
16. Liang, J.H., Ganesh, V., Czarnecki, K., Raman, V.: SAT-based analysis of large real-world feature models is easy. In: Proceedings of the 19th International Software Product Line Conference, SPLC. pp. 91–100 (2015)

17. Marques-Silva, J.: Practical applications of Boolean satisfiability. In: Proceedings of the 9th International Workshop on Discrete Event Systems. pp. 74–80 (2008)
18. Marques-Silva, J., Lynce, I., Malik, S.: Conflict-driven clause learning SAT solvers. In: Biere, A., Heule, M., van Maaren, H., Walsh, T. (eds.) Handbook of Satisfiability, Frontiers in Artificial Intelligence and Applications, vol. 185, chap. 4. IOS Press (2009)
19. Marques-Silva, J.P., Sakallah, K.A.: GRASP - a new search algorithm for satisfiability. In: ICCAD. pp. 220–227 (1996)
20. Mateescu, R.: Treewidth in industrial SAT benchmarks. Tech. Rep. MSR-TR-2011-22, Microsoft Research (February 2011), <http://research.microsoft.com/pubs/145390/MSR-TR-2011-22.pdf>
21. Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B., Troyansky, L.: Determining computational complexity from characteristic phase transitions. *Nature* 400(6740), 133–137 (1999)
22. Newman, M.E., Girvan, M.: Finding and evaluating community structure in networks. *Physical Review E* 69(2), 026113 (2004)
23. Newsham, Z., Ganesh, V., Fischmeister, S., Audemard, G., Simon, L.: Impact of community structure on SAT solver performance. In: Theory and Applications of Satisfiability Testing, 17th International Conference. pp. 252–268 (2014)
24. Robinson, J.A.: A machine-oriented logic based on the resolution principle. *Journal of the ACM (JACM)* 12(1), 23–41 (1965)
25. Samer, M., Szeider, S.: Fixed-parameter tractability. In: Biere, A., Heule, M., van Maaren, H., Walsh, T. (eds.) Handbook of Satisfiability, Frontiers in Artificial Intelligence and Applications, vol. 185, chap. 13. IOS Press (2009)
26. Vizel, Y., Weissenbacher, G., Malik, S.: Boolean satisfiability solvers and their applications in model checking. *Proceedings of the IEEE* 103(11), 2021–2035 (2015)
27. Williams, R., Gomes, C.P., Selman, B.: Backdoors to typical case complexity. In: IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence. pp. 1173–1178 (2003)

A Other Clustering Metrics

In this section we give additional evidence that the PCMs form a broad and interesting class of metrics by showing that several popular clustering metrics other than modularity are PCMs. While these metrics have not previously been used in the context of SAT, they have been widely used in other fields to measure the quality of vertex partitions [1]. Note that when necessary, we have slightly adjusted the definitions to yield values in $[0, 1]$ so that they are graph metrics according to the definition in Section 3.

First we establish some common notation.

Definition 13. *Let P be a path in a weighted graph from s to t . Then*

$$d(s, t) = \min_P \sum_{(u,v) \in P} w(u, v)$$

is the weight of the shortest path from s to t . If there is no path from s to t , let $d(s, t) = \infty$, defining for convenience that $\infty/\infty = 1$.

Definition 14. *For any vertex partition δ , the **community** C_v of v is the equivalence class of v under δ .*

In this section the variable δ always ranges over vertex partitions, which when convenient we view as a set of communities.

Now we can define the *silhouette index*, which measures how separated the communities are from each other [1].

Definition 15. *Let*

$$a(v) = \frac{1}{|C_v|} \sum_{t \in C_v, t \neq v} d(v, t)$$

be the average distance between v and the other vertices in the same community. Let

$$b(v) = \min_{C'_v \neq C_v} \frac{1}{|C'_v|} \sum_{t \in C'_v} d(v, t)$$

*be the average distance between v and the vertices in the closest other community. Then the **silhouette index** of a graph $G = (V, E)$ is*

$$S = \frac{1}{2} \left(1 + \max_{\delta} \frac{1}{|V|} \sum_v \frac{b(v) - a(v)}{\max(a(v), b(v))} \right).$$

Next we define the (external) *conductance*, which compares the weight of the edges spanning communities to the weight of the edges within communities [1].

Definition 16. *If for any $S \subseteq V$ we define*

$$r(S) = \sum_{x \in S} \sum_{y \in V} w(x, y),$$

*then the **conductance** of G is*

$$R = 1 - \max_{\delta} \frac{1}{|\delta|} \sum_{C \in \delta} \frac{\sum_{x \in C} \sum_{y \notin C} w(x, y)}{\min(r(C), r(V \setminus C))}.$$

Another simple metric is *coverage*, which compares the weight of the edges within communities to the total weight of the graph [1].

Definition 17. *The coverage of a graph is*

$$Cov = \max_{\delta} \frac{\sum_{u \in V} \sum_{v \in C_u} w(u, v)}{\sum_{u, v \in V} w(u, v)}.$$

Finally, we define *performance*, which is a sum of two terms: the number of edges within communities, and the number of *missing* edges between communities [1].

Definition 18. *The performance of an unweighted graph $G = (V, E)$ is*

$$Perf = \max_{\delta} \frac{|\{(u, v) \in E : u \in C_v\}| + |\{(u, v) \notin E : u \notin C_v\}|}{n(n-1)}.$$

Note that we only consider the unweighted version of performance, as weighted versions require additional contextual information in the form of reasonable guesses for the weights of missing edges. For all the other metrics above, an unweighted version can be obtained simply by assuming unit weights.

Now we prove that all of these metrics are PCMs. In fact this is for a trivial reason: they all consider a single complete graph to be well-clustered. So an attempt to use them as difficulty metrics in the context of SAT would need to change their definitions, for example by maximizing only over δ with at least some minimum number of communities.

Theorem 2. *Silhouette index, conductance, coverage, and performance are PCMs.*

Proof (sketch). Take $c = 1$, so that $G = K_n$. Consider a partition δ which puts all vertices into the same community. Then there are no vertices outside of that community, so the silhouette index is 1. Similarly, there are no edges between communities, so the conductance and coverage are both 1. Finally, since there are $n(n-1)/2$ edges inside the community, the performance is 1. \square

B Deferred Proofs

Theorem 5. *Suppose that $k \geq 4$, $m = O(n)$, and $c = O(n^\alpha)$ for some $\alpha < \frac{k-2}{4(k-1)}$. Then there is some $\lambda > 0$ so that with high probability, an unsatisfiable formula drawn from $F_k(n, m, c, p)$ does not have a resolution refutation of size $2^{O(n^\lambda)}$.*

Proof. Fix some $\lambda, \epsilon > 0$ which we will require to be sufficiently small later, and define $x(n) = n^{(\lambda-1)/2}$. To apply Theorem 4 we must have $t(n) = \omega(1)$ and

$$t(n) \leq \min\{x(n)n/2, s(n)/z(n)\}$$

where $s(n) = x(n)^{1-1/y(k-1)} n^{1-1/(k-1)}$. So we set $t(n) = x(n)n/2$ and prove that for an appropriate $z(n) = \omega(1)$, this is less than $s(n)/z(n)$ for sufficiently large n . First, note that

$$x(n)^{-1/y} = n^{\frac{-(\lambda-1)}{2y}} = n^{\frac{(1-\lambda)(k+\epsilon)}{4}} = \omega(n)$$

for sufficiently small λ , since $k \geq 4$ and $\epsilon > 0$. Therefore,

$$s(n) = x(n)^{1 - \frac{1}{y(k-1)}} n^{1 - \frac{1}{k-1}} = x(n) \left(x(n)^{-1/y} \right)^{1/(k-1)} n^{1 - \frac{1}{k-1}} \in \omega(x(n)n).$$

Since $t(n) = x(n)n/2$, there is some $z(n)$ in $\omega(1)$ such that $s(n)/z(n) \geq t(n)$ for sufficiently large n . Finally,

$$t(n) = n^{(\lambda-1)/2} n/2 = \Theta \left(n^{(\lambda+1)/2} \right) \subseteq \omega(1)$$

since $\lambda > 0$. Thus we have satisfied all the conditions of Theorem 4. Now observe that

$$\begin{aligned} \frac{1}{2^{7k/2}(1+\epsilon)c^{k-1}} x(n)^{-(k-2-\epsilon)/2} n &= \Omega \left(\frac{1}{n^{\alpha(k-1)}} n^{-(\lambda-1)(k-2-\epsilon)/4} n \right) \\ &= \Omega \left(n^{1 + \frac{(1-\lambda)(k-2-\epsilon)}{4} - \alpha(k-1)} \right) \\ &= \omega \left(n^{1 + \frac{(1-\lambda)(k-2-\epsilon)}{4} - \frac{k-2}{4}} \right) \\ &= \omega \left(n^{1 + \frac{1}{4}[(1-\lambda)(k-2-\epsilon) - (k-2)]} \right) \\ &= \omega(n) \end{aligned}$$

for sufficiently small λ and ϵ . So for sufficiently large n this quantity is larger than $m = O(n)$, and Theorem 4 applies to $F_k(n, m, c, p)$. Therefore with high probability, an unsatisfiable instance from $F_k(n, m, c, p)$ does not have a resolution refutation of size $2^{O(x(n)t(n))} = 2^{O(x(n)^2 n/2)} = 2^{O(n^\lambda)}$. \square

Theorem 6. Suppose $m = O(n)$ and $c = O(n^\alpha)$ for some $\alpha < 1/10$. Then there is some $\lambda > 0$ so that with high probability, an unsatisfiable formula drawn from $F_3(n, m, c, p)$ does not have a resolution refutation of size $2^{O(n^\lambda)}$.

Proof. We proceed along the lines of the previous theorem, except that we set $t(n) = s(n)/z(n) = x(n)^{1-1/2y} n^{1/2}/z(n) = x(n)^{\frac{1-\epsilon}{4}} n^{1/2}/z(n)$ and show that $t(n) \leq x(n) \cdot n/2$ for sufficiently large n .

Fix some $\gamma, \epsilon > 0$ which we will require to be sufficiently small later. Letting $z(n) = n^\gamma$, clearly $z(n) = \omega(1)$. Also define

$$x(n) = \left(n^{-\gamma - \alpha \frac{5-\epsilon}{1-\epsilon}} \right)^{\frac{4}{5-\epsilon}}.$$

Then

$$t(n) = \left(n^{-\gamma - \alpha \frac{5-\epsilon}{1-\epsilon}} \right)^{\frac{4}{5-\epsilon} \cdot \frac{1-\epsilon}{4}} n^{1/2}/z(n) = n^{\frac{1}{2} - \gamma(1 + \frac{1-\epsilon}{5-\epsilon}) - \alpha},$$

which is $\omega(1)$ for sufficiently small γ since $\alpha < 1/10$. Also note that

$$x(n)t(n) = x(n)^{\frac{5-\epsilon}{4}} n^{\frac{1}{2} - \gamma} = n^{\frac{1}{2} - 2\gamma - \alpha \frac{5-\epsilon}{1-\epsilon}} = n^\lambda$$

where $\lambda = \frac{1}{2} - 2\gamma - \alpha \frac{5-\epsilon}{1-\epsilon}$. Again since $\alpha < 1/10$, we have $\lambda > 0$ for sufficiently small γ and ϵ . Similarly,

$$x(n)^{-1/y} = \left(n^{-\gamma - \alpha \frac{5-\epsilon}{1-\epsilon}} \right)^{-\frac{4}{5-\epsilon} \cdot \frac{3+\epsilon}{2}} = n^{2\gamma \frac{3+\epsilon}{5-\epsilon} + 2\alpha \frac{3+\epsilon}{1-\epsilon}} = o(n)$$

for sufficiently small γ and ϵ . Therefore

$$t(n) = x(n) \left(x(n)^{-1/y} \right)^{1/2} n^{1/2} / z(n) = o(x(n)n).$$

So for sufficiently large n we have satisfied the conditions of Theorem 4. Now observe that

$$\begin{aligned} \frac{1}{2^{21/2}(1+\epsilon)c^2} x(n)^{-(1-\epsilon)/2} n &= \Omega \left(\frac{1}{n^{2\alpha}} \left(n^{-\gamma-\alpha \frac{5-\epsilon}{1-\epsilon}} \right)^{-\frac{4}{5-\epsilon} \cdot \frac{1-\epsilon}{2}} n \right) \\ &= \Omega \left(n^{1+2\gamma \frac{1-\epsilon}{5-\epsilon}} \right) \\ &= \omega(n) \end{aligned}$$

since $\gamma > 0$ and we may take $\epsilon < 1$. So for sufficiently large n this quantity is larger than $m = O(n)$, and Theorem 4 applies to $F_3(n, m, c, p)$. Therefore with high probability, an unsatisfiable instance from $F_3(n, m, c, p)$ does not have a resolution refutation of size $2^{O(x(n)t(n))} = 2^{O(n^\lambda)}$. \square

Theorem 8. *Let $\tilde{F}_k(n, m, c, p)$ be the community attachment model modified so that any clause that is not within a single community counts as a bridge clause. Then if $m = O(n)$ and $c = O(n^\alpha)$ for any $\alpha < 1/10$, the runtime of CDCL on an unsatisfiable formula φ from $\tilde{F}_k(n, m, c, p)$ is exponential with high probability.*

Proof. Modify Algorithm 1 to use the new definition of bridge clause, store in b the new bridge clause probability $1 - h$, and sample from $\tilde{F}_k(n, m, c, p)$ on line 15.

Now we check that each lemma is still true. For Lemma 1, observe that removing the constraint that every variable in a bridge clause must come from a different community cannot decrease the probability that a random clause is a bridge clause. So our new value of b is at least as large as the old, and therefore h/b is still at most 1. Also $\tilde{F}_k(n, m, c, p)$ is a distribution over k -CNF formulas with n variables and m clauses, so Lemma 1 holds.

For Lemma 2, when the modified Algorithm 1 returns from line 16 the formula ψ is drawn from $\tilde{F}_k(n, m, c, p)$. So in this case $\bar{F}_k(n, m, c, p; m')$ is trivially identical to $\tilde{F}_k(n, m, c, p)$, and we need only consider the case when the algorithm returns from line 14. As above each clause of ϕ is added independently, so we need only calculate the probability that an added clause is within a community. Proceeding exactly as in Lemma 2, we obtain $\Pr[C_{\text{added}}|A] = \Pr[C_{\text{comm}}] = h$ and $\Pr[C_{\text{added}}|\bar{A}] = (h/b) \Pr[C_{\text{bridge}}] = h$ (since we changed b to be the probability of getting a bridge clause under the new definition). So

$$\Pr[C_{\text{added}}] = \Pr[C_{\text{added}}|A] \Pr[A] + \Pr[C_{\text{added}}|\bar{A}] \Pr[\bar{A}] = hp + h(1-p) = h,$$

and therefore

$$\Pr[C_{\text{comm}}|C_{\text{added}}] = \Pr[A|C_{\text{added}}] = \frac{\Pr[C_{\text{added}}|A] \Pr[A]}{\Pr[C_{\text{added}}]} = \frac{hp}{h} = p$$

as before.

For Lemma 4, note that the probability that C is a bridge clause is the new value of b . So as before the probability that C is added to ψ is $p \cdot h + (1 - p) \cdot (h/b) \cdot b = h$. The rest of the computation then proceeds without change.

Finally, the argument for Lemma 5 goes through with no changes. So subformula-inherited properties of random k -SAT instances are passed on to instances of $\tilde{F}_k(n, m, c, p)$ with high probability. Therefore the Beame–Pitassi argument in Section 4.4 and the CDCL runtime bound in Section 4.5 hold for $\tilde{F}_k(n, m, c, p)$ with no further modifications needed. \square